REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE December 1995	3. REPORT TYPE AND DATES COVERED final technical 09/01/92--08/31/95

4. TITLE AND SUBTITLE
Neural Models of Spatial Orientation
in Novel Environments

5. FUNDING NUMBERS
F49620-92-J-0499
2313 - AS
6110 2F

6. AUTHOR(S)
Professor Stephen Grossberg, Principal Investigator

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Boston University
Center for Adaptive Systems
and
Department of Cognitive and Neural Systems
Boston, MA 02215

8. PERFORMING ORGANIZATION REPORT NUMBER

AFOSR-TR-95
0789

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Office of Scientific Research
Chemistry & Life Sciences Directorate
Building 410
Bolling AFB, DC 20332-6448

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

12b. DISTRIBUTION CODE
JAN 0 4 1996
F

13. ABSTRACT (Maximum 200 words)

Completed work on this grant analyzes problems of spatial navigation, 3-D vision, visual search, spatial mapping, planning and action, all of which utilize spatial representations to control adaptive behaviors in real time. Highlights include contributions towards solving the classical figure-ground problem for biological vision, the problem of self-organizing body-centered spatial representations for movement planning and spatial orientation, and the problem of carrying out fast visual search for targets among multiple distractors. Ongoing research includes projects concerning how a rapidly moving agent can self-organize spatial representations, use these representations for real-time movement planning, and transform spatial movement plans into appropriate motor commands for movement control and real-time navigation in a rapidly changing, cluttered environment. Specific projects include retinal image processing, formation of egocentric maps of heading direction from optic flow, detection of independently moving objects from optic flow, integration of egocentric and allocentric representations for autonomous navigation, and investigation of spatial reference frames for real-time flexible speech articulator control.

14. SUBJECT TERMS

15. NUMBER OF PAGES
43 pages

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
unclassified			

NSN 7540-01-280-5500

DTIC QUALITY INSPECTED 1

1996010 3 112

Final Technical Report
September 1, 1992—August 31, 1995

NEURAL MODELS OF SPATIAL ORIENTATION IN NOVEL ENVIRONMENTS

Contract AFOSR F49620-92-J-0449

Professor Stephen Grossberg, PI
Boston University
Department of Cognitive and Neural Systems
and
Center for Adaptive Systems
677 Beacon Street
Boston, MA 02215
(617) 353-7857 (phone)
(617) 353-7755 (fax)

TABLE OF CONTENTS

PUBLICATIONS PARTIALLY SUPPORTED BY
THE AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
CONTRACT AFOSR F49620-92-J-0499
SEPTEMBER 1, 1992—AUGUST 31, 1995
CENTER FOR ADAPTIVE SYSTEMS
AND
DEPARTMENT OF COGNITIVE AND NEURAL SYSTEMS
BOSTON UNIVERSITY

1. Bradski, G. and Grossberg, S. (1994). A neural architecture for 3-D object recognition from multiple 2-D views. In **Proceedings of the world congress on neural networks**, San Diego, **IV**, 211–219. Hillsdale, NJ: Erlbaum Associates. (*+@)

2. Bradski, G. and Grossberg, S. (1994). Recognition of 3-D objects from multiple 2-D views by a self-organizing neural architecture. In V. Cherkassky, J.H. Friedman, and H. Wechsler (Eds.), **From statistics to neural networks: Theory and pattern recognition**. New York: Springer-Verlag, pp. 349–375. (*+@)

3. Bradski, G. and Grossberg, S. (1995). VIEWNET: A neural architecture for learning to recognize 3-D objects from multiple 2-D views. In **Proceedings of the conference on intelligent robots and computer vision XIII: Algorithms and computer vision**. Bellingham, WA: SPIE, in press. (*+@)

4. Bradski, G. and Grossberg, S. (1995). Fast learning VIEWNET architectures for recognizing 3-D objects from multiple 2-D views. **Technical Report CAS/CNS-TR-93-053**, Boston University. *Neural Networks* (Special Issue on Automatic Target Recognition), **8**, in press. (*+)

5. Bullock, D., Cisek, P.E., and Grossberg, S. (1995). Cortical networks for control of voluntary arm movements under variable force conditions. **Technical Report CAS/CNS-TR-95-019**, Boston University. Submitted for publication. (+@)

6. Bullock, D., Contreras-Vidal, J.L. and Grossberg, S. (1993). Cerebellar learning in an opponent motor controller for adaptive load compensation and synergy formation. In **Proceedings of the world congress on neural networks**, Portland, **IV**, 481–486. Hillsdale, NJ: Erlbaum Associates. (*@)

7. Bullock, D., Greve, D., Grossberg, S., and Guenther, F.H. (1993). A self-organizing neural network for learning a body-centered invariant representation of 3-D target position. In **Proceedings of the world congress on neural networks**, Portland, **II**, 405–408. Hillsdale, NJ: Erlbaum Associates. (+)

8. Bullock, D., Greve, D., Grossberg, S., and Guenther, F.H. (1993). A self-organizing neural network for learning a body-centered invariant representation of 3-D target position. In S. Gielen and B. Kappen (Eds.), **Proceedings of the international conference on artificial neural networks (ICANN'93)**. London: Springer-Verlag, pp. 90–95. (+)

9. Bullock, D., Grossberg, S., and Guenther, F.H. (1993). A self-organizing neural model of motor equivalent reaching and tool use by a multijoint arm. *Journal of Cognitive Neuroscience*, **5**, 408–435. (+@)

10. Bullock, D., Grossberg, S., and Guenther, F.H. (1995). Neural network modeling of sensory-motor control in animals. In H. Zelaznik (Ed.), **Advances in motor learning and control**. Champaign, IL: Human Kinetics Press, in press. (+@)

11. Bullock, D., Grossberg, S., and Mannes, C. (1993). A neural network model for cursive script production. *Biological Cybernetics*, **70**, 15-28. (*+@)

12. Cameron, S., Grossberg, S., and Guenther, F.H. (1995). A self-organizing heading and depth detection network. In **Proceedings of the world congress on neural networks**, San Diego, **I**, 3–7. Hillsdale, NJ: Erlbaum Associates. (*@)

13. Carpenter, G.A. and Grossberg, S. (1995). Adaptive resonance theory. In J.D. Irwin (Ed.), **The industrial electronics handbook**. Boca Raton: CRC Press, in press. (*+@)

14. Carpenter, G.A., Grossberg, S., and Lesher, G. (1993). The what-and-where filter: A spatial mapping neural network for object recognition and image understanding. **Technical Report CAS/CNS-TR-93-043**, Boston University. Submitted for publication. (*#+@)

15. Chey, J., Cisek, P., Gaudiano, P., and Wood, R. (1994). Are learned biases in *Aplysia* head-waving due to specific or nonspecific mechanisms? *Society for Neuroscience Abstracts*, **20**(2), 1073. (*@)

16. Chey J., Cisek P., Gaudiano P., and Wood R.J. (1994). A neural model of biased oscillations in Aplysia head-waving behavior. **Technical Report CAS/CNS-TR-94-031**, Boston University.

17. Chey, J., Grossberg, S., and Mingolla, E. (1995). Neural dynamics of motion processing and speed discrimination. **Technical Report CAS/CNS-TR-94-030**, Boston University. *Vision Research*, in press. (*+@)

18. Dorman, C. and Gaudiano, P. (1995). Motivation. **Technical Report CAS/CNS-TR-94-020**, Boston University. In M.A. Arbib (Ed.), **The handbook of brain theory and neural networks**. Cambridge, MA: MIT Press, in press.

19. Fiala, J. and Guenther, F.H. (1994). Book Review: **Handbook of intelligent control: Neural, fuzzy, and adaptive approaches**, D.A. White and D.A. Sofge (Eds.). *Neural Networks*, **7**, 851–852.

20. Francis, G. and Grossberg, S. (1995). Cortical dynamics of boundary segmentation and reset: Persistence of residual traces and orientational aftereffects. **Technical Report CAS/CNS-TR-95-002**, Boston University. Submitted for publication. (*@)

21. Gaudiano, P. (1993). An unsupervised error-based neural network model for the generation and control of movement trajectories. In **Proceedings of the AAAS Science Innovation Conference, 93-21S**, 108–109.

22. Gaudiano, P. (1994). Book Review: **Motivation, emotion, and goal direction in neural networks** by Daniel S. Levine and Samuel J. Leven. *Neural Networks*, **7**, 213–214.

23. Gaudiano, P. (1994). Simulations of X and Y retinal ganglion cell behavior with a nonlinear push-pull model of spatiotemporal retinal processing. *Vision Research*, **34**, 1767–1784.

24. Gaudiano, P., Guenther, F.H., and Zalama, E. (1995). The neural dynamics approach to sensory-motor control: Overview and recent applications in mobile robot control and speech production. Invited article to appear in P.P. van der Smagt (Ed.), **Progress in neural networks** (special volume on Neural Systems for Robotics). Norwood, NJ: Ablex Publishing Corporation, in press.

25. Gaudiano P., Linsay P.S., Przybyszewski A.W., and Wilson C.M. (1995). Coupled nonlinear oscillators as a model of retinal ganglion cells. *Abstracts of the 1995 Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO)*, #4280.

26. Gaudiano, P. and Mejía-Monasterio, N.A. (1994). An analysis of psychophysical and physiological adaptation with a dynamic photoreceptor model. In **Proceedings of the 17th European conference on visual perception**, Eindhoven, The Netherlands.

27. Gaudiano, P., Olson, S., Tal, D., and Fischl, B. (1993). A neural network model of dynamic receptive field reorganization. *Society for Neuroscience Abstracts*, **19**, #328.19.

28. Gaudiano, P. and Pessoa, L. (1993). Dynamics of temporal sensitivity with the push-pull shunting model of retinal processing. In **Proceedings of the 16th European conference on visual perception**, Edinburgh, Scotland. Abstract appears in *Perception*, **22**, 142.

29. Gaudiano, P. and Pessoa, L. (1994). Bloch's law, critical duration, and temporal sensitivity with the push-pull model of retinal function. *Investigative Ophthalmology and Visual Science Supplement*, Abstract #2191.

30. Gaudiano, P. and Rosen, D.B. (1994). Book Review: **Introduction to the theory of neural computation** by John Hertz, Anders Krogh, and Richard G. Palmer. *Neural Networks*, **7**, 1325–1327.

31. Gaudiano, P., Surmeli, D., and Wilson, F. (1994). Gated dipoles for operant conditioning. *Neural Networks*, **7**, 405–407.

32. Gaudiano, P., Zalama, E., and López-Coronado, J. (1995). An unsupervised neural network for low-level control of a wheeled mobile robot. **Technical Report CAS/CNS-TR-94-019**, Boston University. *IEEE Transactions on Systems, Man, and Cybernetics*, in press.

33. Gove, A., Grossberg, S., and Mingolla, E. (1995). Brightness perception, illusory contours, and corticogeniculate feedback. **Technical Report CAS/CNS-TR-94-033**, Boston University. *Visual Neuroscience*, **12**, 1027–1052. (*@)

34. Grossberg, S. (1993). Boundary, brightness, and depth interactions during preattentive representation and attentive recognition of figure and ground. *Italian Journal of Psychology*, **XX**, 771–804. Invited article in celebration of the 80th birthday of Professor Gaetano Kanizsa and the 20th anniversary of the Italian Journal of Psychology. (*@)

35. Grossberg, S. (1994). 3-D vision and figure-ground separation by visual *Perception and Psychophysics*, **55**, 48–120. (*@)

36. Grossberg, S. (1995). Figure-ground separation. **Technical Report CAS/CNS-TR-93-062**, Boston University. In M.A. Arbib (Ed.), **Handbook of brain theory and neural networks**. Cambridge, MA: MIT Press, in press. (*@)

37. Grossberg, S. (1995). Letter to the Editor: Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, in press. (*@)

38. Grossberg, S. and Grunewald, A. (1994). Spatial pooling and perceptual framing by synchronizing cortical dynamics. In M. Marinaro and P.G. Morasso (Eds.), **Proceedings of the international conference on artificial neural networks (ICANN-94)**. New York: Springer-Verlag, pp. 10–15. (*@)

39. Grossberg, S. and Grunewald, A. (1994). Synchronized neural activities: A mechanism for perceptual framing. In **Proceedings of the world congress on neural networks**, San Diego, **IV**, 655–660. Hillsdale, NJ: Erlbaum Associates. (*@)

40. Grossberg, S. and Grunewald, A. (1994). Cortical synchronization and perceptual framing. **Technical Report CAS/CNS-TR-94-025**, Boston University. Submitted for publication. (*@)

41. Grossberg, S. and Grunewald, A. (1995). Temporal dynamics of binocular disparity processing by corticogeniculate interactions. **Technical Report CAS/CNS-TR-95-021**, Boston University. Submitted for publication. (*@)

42. Grossberg, S., Mingolla, E., and Ross, W.D. (1994). A neural theory of attentive visual search: Interactions of visual, spatial, and object representations. *Psychological Review*, **101**, 470–489. (*+@)

43. Grossberg, S., Mingolla, E., and Williamson, J. (1995). Synthetic aperture radar processing by a multiple scale neural system for boundary and surface representation. **Technical Report CAS/CNS-TR-94-001**, Boston University. *Neural Networks* (Special Issue on Automatic Target Recognition), **8**, in press. (++@)

44. Grossberg, S., Mingolla, E., and Williamson, J. (1995). A multiple scale neural system for boundary and surface representation of SAR data. In **Proceedings of the IEEE workshop on neural networks for signal processing**. New York: IEEE Publishing Services, in press. (*+@)

45. Grossberg, S. and Olson, S.J. (1994). Rules for the cortical map of ocular dominance and orientation columns. *Neural Networks*, **7**, 883–894. (*#+@)

46. Grossberg, S. and Olson, S.J. (1994). Rules for the cortical map of ocular dominance and orientation columns. In **Proceedings of the world congress on neural networks**, San Diego, **IV**, 284–289. Hillsdale, NJ: Erlbaum Associates. (*#+@)

47. Grossberg, S., Pribe, C., and Cohen, M.A. (1994). Neural control of interlimb oscillations, I: Human bimanual coordination. **Technical Report CAS/CNS-TR-94-021**, Boston University. Submitted for publication. (*+@)

48. Grunewald, A. and Grossberg, S. (1994). Binding of object representations by synchronous cortical dynamics explains temporal order and spatial pooling data. In A. Ram and K. Eiselt (Eds.), **Proceedings of the sixteenth annual conference of the Cognitive Science Society**. Hillsdale, NJ: Erlbaum Associates, pp. 387–391. (*@)

49. Guenther, F.H. (1993). A self-organizing neural model for motor equivalent phoneme production. In **Proceedings of the world congress on neural networks**, Portland, **III**, 6–9. Hillsdale, NJ: Erlbaum Associates.

50. Guenther, F.H. (1993). A self-organizing neural model for motor equivalent phoneme production. In S. Gielen and B. Kappen (Eds.), **Proceedings of the international conference on artificial neural networks (ICANN'93)**. London: Springer-Verlag, pp. 71–74.

51. Guenther, F.H. (1993). Sensorimotor transformations in a neural model of motor equivalent speaking. *Society for Neuroscience Abstracts*, **19**(1), p. 553.

52. Guenther, F.H. (1994). A neural network model of speech acquisition and motor equivalent speech production. *Biological Cybernetics*, **72**, 43–53.

53. Guenther, F.H. (1994). Skill acquisition, coarticulation, and rate effects in a neural network model of speech production. **Program of the 127th meeting of the Acoustical Society of America**. *Journal of the Acoustical Society of America*, **95**(5), p. 2924.

54. Guenther, F.H. (1995). Speech sound acquisition, coarticulation, and rate effects in a neural network model of speech production. **Technical Report CAS/CNS-TR-94-012**, Boston University. *Psychological Review*, **102**, 594–621.

55. Guenther, F.H. (1995). A modeling framework for speech motor development and kinematic articulator control. In K. Elenius and P. Branderud (Eds.), **Proceedings of the XIIIth international congress of phonetic sciences** (Stockholm, Sweden, 13–19 August 1995). Stockholm: KTH and Stockholm University, Volume 2, pp. 92–99.

56. Guenther, F.H., Bullock, D., Greve, D., and Grossberg, S. (1994). Neural representations for sensory-motor control, III: Learning a body-centered representation of 3-D target position. *Journal of Cognitive Neuroscience*, **6**, 341–358. (+@)

57. Guenther, F.H. and Johnson, D. (1995). A computational model using formant space planning of articulator movements for vowel production. *Journal of the Acoustical Society of America*, **97**(5), p. 3402.

58. Guenther, F.H. and Micci-Barreca, D. (1995). Efficient curved reaches resulting from kinematic biases in the DIRECT model. In **Proceedings of the IEEE international conference on systems, man, and cybernetics**, Vancouver, Canada.

59. Guenther, F.H. and Meyers, C. (1995). Book Review: **An introduction to the modeling of neural networks** by P. Peretto. *Neural Networks*, in press. Also to appear in *SIAM Review*.

60. Johnson, D. and Guenther, F.H. (1995). Acoustic space movement planning in a neural model of motor equivalent vowel production. In **Proceedings of the world congress on neural networks**, Washington, **1**, 481–485. Hillsdale, NJ: Erlbaum Associates.

61. Linsay P.S., Wilson C.M., Przybyszewski, A.W., and Gaudiano P. (1995). Modeling of neural signal processing with hardware nonlinear oscillators. *Abstracts of the Third SIAM Conference on Applications of Dynamical Systems*, Snowbird, Utah.

62. Mejia-Monasterio, N.A. and Gaudiano, P. (1994). Dynamic light adaptation to aperiodic stimuli with the push-pull model of retinal function. *Investigative Ophthalmology and Visual Science Supplement*, Abstract #2679.

63. Mejía-Monasterio, N.A. and Gaudiano P. (1995). A calcium-dependent mechanism improves light adaptation in a simulated photoreceptor. *Abstracts of the 1995 Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO)*, #3041.

64. Micci-Barreca, D. and Guenther, F.H. (1995). Efficient trajectory formation using a learned approximate pseudoinverse in the DIRECT model of reaching. In **Proceedings of the world congress on neural networks**, Washington, **1**, 388–392. Hillsdale, NJ: Erlbaum Associates.

65. Muñiz F., Zalama E., Gaudiano P., and López Coronado J. (1995). Neural controller for a mobile robot in a nonstationary environment. In **Proceedings of the second IFAC conference on intelligent autonomous vehicles**, Helsinki, Finland, June 1995, pp. 279–284.

66. Olson S., Fischl B., and Gaudiano P. (1995). A model of rapid and long-term reorganization in adult visual cortex. *Abstracts of the 1995 Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO)*, #1786.

67. Pedini, D. and Gaudiano, P. (1994). A collection of ART-family graphical simulators. Submitted for publication. (+)

68. Pribe, C., Grossberg, S., and Cohen, M.A. (1994). Neural control of interlimb oscillations, II: Biped and quadruped gaits and bifurcations. **Technical Report CAS/CNS-TR-94-022**, Boston University. Submitted for publication. (*+@)

69. Przybyszewski, A.W., Gaudiano P., and Linsay P.S. (1995). Coupled oscillators in the retina as a model of brain function. *Abstracts of the Third SIAM Conference on Applications of Dynamical Systems*, Snowbird, Utah.

70. Przybyszewski A.W., Gaudiano P., and Sucher N.J. (1995). APB modification of spatiotemporal processes in the cat retina: Comparing experiment and simulation. *Abstracts of the 1995 Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO)*, #1321.

71. Przybyszewski A.W., Gaudiano P., van Wezel R., and van de Grind W.A. (1995). Spatial asymmetry in cat retinal ganglion cells light responses. *Society for Neuroscience Abstracts*, #211.6, 509.

72. Przybyszewski, A.W., Linsay P.S., Gaudiano P., and Wilson C.M. (1995). Hardware coupled nonlinear oscillators as a model of retina. In **Proceedings of the third experimental chaos conference**, Edinburgh, Scotland, August 1995. Also appears as **Technical Report CAS/CNS-TR-95-028**, Boston University.

73. Zalama, E., Gaudiano, P., and Lopéz Coronado, J. (1994). A real-time, unsupervised neural network for the control of a mobile robot. **Technical Report CAS/CNS-TR-93-056**, Boston University. In **Proceedings of the IEEE world congress on computational intelligence**, Orlando, Florida.

74. Zalama, E., Gaudiano, P., and Lopéz Coronado, J. (1995). A real-time, unsupervised neural network for the low-level control of a mobile robot in a nonstationary environment. *Neural Networks*, **8**, 103–123.

75. Zalama E., Gaudiano P., and López Coronado J. (1995). Obstacle avoidance by means of an operant conditioning model (oral presentation). In José Mira and Francisco Sandoval (Eds.), **Lecture notes in computer science: From natural to artificial neural computation**, International Workshop on Artificial Neural Networks, Malaga-Torremolinos, Spain, June 1995, pp. 471–477.

RESEARCH SUMMARIES

1. Designing a Self-Organizing Network that Calculates Translational Heading [12]

Navigation requires heading perception. For a robot to maintain a constant heading, it is insufficient to simply set the wheels at a given specific angle and assume the resulting trajectory is correct. An autonomous vehicle on a side-hill, for example, may have to turn its wheels just to maintain a constant heading due to tire slippage. Visual feedback is a more reliable indicator of heading. Algorithms exist that allow explicit calculation of heading from optic flow, but none are self-organizing and each requires having explicit knowledge about the geometry of the sensor, the focal length of the optics, etc. Such parameters must be recalculated for each variety of robot and if they change over time, the robot will begin to malfunction.

This project investigates a self-organizing solution for determining heading from optic flow. Such a solution is generic, robust, and automatically adapts to a wide variety of robotic systems. A network-like solution has been developed based on the subspace algorithm introduced by Heeger and Jepson (1992). The subspace algorithm calculates the translational component of 3D motion without explicit knowledge of depth or rotational velocity. Its stability, generality, and efficiency place it among the best techniques available. Lappe and Rauschecker (1993) first noticed the potential of expressing the algorithm in terms of a neural model, but their model is not self-organizing.

Its weaknesses are:
(1) It only handles forward translations in the Y-plane;
(2) The weights need to be calculated off-line;
(3) It assumes heading direction cells in the heading map are predetermined.

This work has made the following improvements:
(1) The network is self-organizing for both the heading field and the weight field;
(2) The network determines heading for all directions in the sphere, not just 2-dimensional forward motion.

A schematic diagram of the heading network is shown in Figure 1.

2. Building a Range Map from Optic Flow Information [12]

A second requirement for navigation is knowledge of the 3D spatial arrangement of objects in the environment. Information about the depth of objects is encoded in optic flow but somewhat difficult to extract when arbitrary motions of the observer are allowed. This project investigates how to use information from the heading network to build a range map of the visual scene.

Since two unknowns, translational velocity and inverse depth, are multiplied by each other in the optic flow equation, optic flow cannot give information about absolute speed of the navigator or absolute depth of objects. However, a self-organizing network has been developed that yields relative inverse depth information. When travelling parallel the optic axis, the relative depth information becomes equivalent to a time-to-collision measurement.

The network, shown in Figure 2, uses extraretinal information regarding eye rotational velocities and heading direction from the previously describe network to self-organize into a range map. Future work will involve placing the eye-centered range information into a body-centered representation for navigation.

3. Neural Controller for a Mobile Robot ([24], [32], [65], [73]–[75]) Since the earliest stages of this project, Professor Gaudiano has been involved with the development

Figure 1. Self organizing network for extracting eye-centered ego-motion.


Figure 2. Self-organizing range network.

of a neural network for the low-level control of a mobile robot. This line of work has been extended during the course of the project year, and is the subject of further research. A first manuscript describing the neural controller will appear shortly in the international journal *Neural Networks*. A second article has been accepted (pending some revisions) for publication in a special issue of the *IEEE Transactions on Systems, Man, and Cybernetics*. The latter article provides more rigorous tests and analyses of noise resistance in the neural controller; it proves asymptotic stability of the neural control scheme; and it illustrates some initial attempts to implement the neural controller using the commercially available robot *ROBUTER*. The results obtained to date demonstrate that the proposed neural controller offers a viable alternative to more traditional controllers, especially when the robot must operate in an unknown or nonstationary environment.

We have also begun formal development of a neural architecture for navigation and obstacle avoidance, to be used in conjunction with the low-level controller. The architecture, which had been sketched in an earlier progress report, combines information about target and obstacles into a spatial map. The target is represented as a positive Gaussian distribution of activity, while each obstacle is represented as a negative Gaussian. Combination of all Gaussians into a spatial map generates an activity surface similar to the functions used in so-called "potential field" approaches to navigation, which typically utilize gradient descent procedures to locate the maximum of the potential field corresponding to the given configuration of target and obstacles. Unlike potential field approaches, the approach proposed here does not make use of gradient descent: instead, the competition in the spatial map makes it possible to extract a time-varying peak whose position reflects the moment-by-moment direction that the robot must follow in order to reach the target while navigating around obstacles.

In the most recent progress report, a diagram had been included to illustrate how such a scheme might be implemented within the context of a *conditioning circuit*, a neural network model proposed by Grossberg to explain how humans and animals can learn to exhibit appropriate avoidance or approach behavior in response to a given set of external sensory cues, while also taking into account the human's (or animal's) internal needs. In that proposal, it was suggested that each item in the environment that represents an obstacle or target, is stored in memory along with its position in space. However, it was suggested that each object's position would be stored in allocentric (i.e., relative to a fixed external frame) Cartesian coordinates. We have found that a better performance is obtained by utilizing instead an egocentric polar map, in which the position of the target and obstacles are continuously updated as the robot moves in space. With this scheme we have been able to demonstrate robust obstacle avoidance in the presence of moving or stationary target and obstacles, for a variety of configurations of obstacles. We are currently extending this work to include modulation of the Gaussians as a function of internal needs, for example, as a function of the robot's need to recharge its batteries. We are developing schemes to avoid getting caught in "local minima", for instance when the robot is surrounded by obstacles, while the only possible path out of the obstacle area is in a direction opposite that of the target. We expect to present our preliminary results on this work at a relevant conference, and to complete a full-length publication describing this work sometime in 1995.

4. Neural Networks for Learning a Body-Centered Representation of 3-D Target Position [56]

This work, which appeared in the *Journal of Cognitive Neuroscience*, represents the third article in a series aimed at modeling how the brain autonomously learns spatial representations capable of controlling flexible goal-oriented movements. In it, Grossberg and Guenther, working with Dan Bullock and Doug Greve, designed a neural network to model how the brain may autonomously learn a body-centered representation of 3-D target position by combining information about retinal target position, eye position, and head position in real time. Such a body-centered spatial representation enables accurate movement commands to the

limbs to be generated despite changes in the spatial relationships between the eyes, head, body, and limbs through time. The representation is a vector representation—otherwise known as a parcellated distributed representation—of target vergence with respect to the two eyes, and of the horizontal and vertical spherical angles of the target with respect to a cyclopean egocenter. A similar representation has been reported in the caudal midbrain and medulla of the frog, as well as in psychophysical movement studies in humans. A head-centered vector representation of this type is generated by two stages of opponent processing that combine corollary discharges of outflow movement signals to the two eyes. This head-centered vector representation interacts with representations of neck movement commands to generate a body-centered estimate of target position. The contributions of the neck command signals to this vector representation are learned during head movements made while the gaze remains fixed on a foveated target. An initial estimate is stored and offset of a gating signal prevents the stored estimate from being reset during a gaze-maintaining head movement. As the head moves, new estimates are generated and compared with the stored estimate. If the estimates are unequal, the comparison generates non-zero difference vectors, which act as error signals to drive the learning process.

5. 3-D Vision and Figure-Ground Separation by Visual Cortex [35]

In this work, which appeared in *Perception and Psychophysics*, Grossberg further developed a neural network theory of 3-D vision, called FACADE Theory. The theory proposes a solution of the classical figure-ground problem for biological vision. It does so by suggesting how boundary representations and surface representations are formed within a Boundary Contour System (BCS) and a Feature Contour System (FCS). The BCS and FCS interact reciprocally to form 3-D boundary and surface representations that are mutually consistent. Their interactions generate 3-D percepts wherein occluding and occluded object parts are separated, completed, and grouped. The theory clarifies how preattentive processes of 3-D perception and figure-ground separation interact reciprocally with attentive processes of spatial localization, object recognition, and visual search. A new theory of stereopsis is proposed that predicts how cells sensitive to multiple spatial frequencies, disparities, and orientations are combined by context-sensitive filtering, competition, and cooperation to form coherent BCS boundary segmentations. Several factors contribute to figure-ground pop-out, including: boundary contrast between spatially contiguous boundaries, whether due to scenic differences in luminance, color, spatial frequency, or disparity; partially ordered interactions from larger spatial scales and disparities to smaller scales and disparities; and surface filling-in restricted to regions surrounded by a connected boundary. Phenomena such as 3-D pop-out from a 2-D picture, DaVinci stereopsis, 3-D neon color spreading, completion of partially occluded objects, and figure-ground reversals are analysed. The BCS and FCS subsystems model aspects of how the two parvocellular cortical processing streams that join the Lateral Geniculate Nucleus to prestriate cortical area V4 interact to generate a multiplexed representation of Form-And-Color-And-DEpth, or FACADE, within area V4. Area V4 is suggested to support figure-ground separation and to interact with cortical mechanisms of spatial attention, attentive object learning, and visual search. Adaptive Resonance Theory (ART) mechanisms model aspects of how prestriate visual cortex interacts reciprocally with a visual object recognition system in inferotemporal cortex (IT) for purposes of attentive object learning and categorization. Object attention mechanisms of the What cortical processing stream through IT cortex are distinguished from spatial attention mechanisms of the Where cortical processing stream through parietal cortex. Parvocellular BCS and FCS signals interact with the model What steam. Parvocellular FCS and magnocellular Motion BCS signals interact with the model Where stream. Reciprocal interactions between these visual, What, and Where mechanisms are used to discuss data about visual search and saccadic eye movements, including fast search of conjunctive targets, search of 3-D surfaces, selective search of like-colored targets, attentive tracking of multi-element groupings, and recursive search of simultaneously presented targets. These interactions shed new light on

10

how spatial and recognition mechanisms interact to identify targets moving in space.

6. A Neural Theory of Attentive Visual Search: Interactions of Visual, Spatial, and Object Representations [42]

Visual search by humans for targets distributed in space among distractors is remarkably efficient. In this work, which appeared in *Psychological Review*, Grossberg, Mingolla, and Ross show how visual search data can be given a unified quantitative explanation by a model of how spatial maps in the parietal cortex and object recognition categories in the inferotemporal cortex deploy attentional resources as they reciprocally interact with visual representations in the prestriate cortex. The model visual representations are organized into multiple boundary and surface representations. Visual search in the model is initiated by organizing multiple items that lie within a given boundary or surface representation into a candidate search grouping. These items are matched with object recognition categories to test for matches or mismatches. Mismatches can trigger deeper searches and recursive selection of new groupings until a target object is identified. This search model is algorithmically specified to quantitatively simulate search data using a single set of parameters, as well as to qualitatively explain a still larger data base, including data of Aks and Enns (1992), Bravo and Blake (1990), Egeth, Virzi, and Garbart (1984), Cohen and Ivry (1991), Enns and Rensink (1990), He and Nakayama (1992), Humphreys, Quinlan, and Riddoch (1989), Mordkoff, Yantis, and Egeth (1990), Nakayama and Silverman (1986), Treisman and Gelade (1980), Treisman and Sato (1990), Wolfe, Cave, and Franzel (1989), and Wolfe and Friedman-Hill (1992). The model hereby provides an alternative to recent variations on the Feature Integration and Guided Search models, and grounds the analysis of visual search in neural models of preattentive vision, attentive object learning and categorization, and attentive spatial localization and orientation.

7. The What-and-Where Filter: A Spatial Mapping Neural Network for Object Recognition and Image Understanding [14]

In this work, Carpenter, Grossberg, and Greg Lesher have modeled a What-and-Where filter that forms part of a neural network architecture for spatial mapping, object recognition, and image understanding. The Where filter responds to an image figure that has been separated from its background. It generates a spatial map whose cell activations simultaneously represent the position, orientation, and size of the figure (where it is). This spatial map may be used to direct spatially localized attention to these image features. A multiscale array of oriented detectors, followed by competitive interactions between position, orientation, and size scales, is used to define the Where filter. The Where filter may be used to transform the image figure into an invariant representation that is insensitive to the figure's original position, orientation, and size. This invariant figural representation forms part of a system devoted to attentive object learning and recognition (what it is). The Where spatial map of all the figures in an image, taken together with the invariant recognition categories that identify these figures, can be used to learn multidimensional representations of objects and their spatial relationships for purposes of image understanding. The What-and-Where filter is inspired by neurobiological data showing that a Where processing stream in the cerebral cortex is used for attentive spatial localization and orientation, whereas a What processing stream is used for attentive object learning and recognition.

8. A Collection of ART-Family Graphical Simulators [67]

The Adaptive Resonance Theory (ART) architecture, first proposed by Grossberg, is a self-organizing neural network for stable pattern categorization in response to arbitrary input sequences. Since its original formulation, several versions of ART have been proposed, each designed to handle a particular task or input format. Recent ART architectures have been developed by Carpenter and Grossberg working with CNS graduate students. Some of

these models have been designed to work in a supervised fashion, offering a viable alternative to supervised neural networks such as backpropagation. Perhaps the best-known variant of ART is ART 2, an unsupervised neural network that handles analog inputs. In collaboration with a graduate student, Professor Gaudiano has developed a series of simulators for some of the ART-family neural architectures, namely, ART 2, ART2-A, Fuzzy ART, and Fuzzy ARTMAP. These simulators, which will soon become available on the public domain, are written in C++ and utilize the Tcl/Tk language for the graphical user interface. The software was implemented on UNIX and NeXT workstations, and should be easily ported to other platforms. It is expected that this software package will be utilized extensively both as a pedagogical tool and as a tool for research around the world.

9. Reactive Obstacle Avoidance [75]

During the final year of the project, Professor Gaudiano continued to extend his research on adaptive control of autonomous mobile robots. The work has been extended to include modules for navigation and obstacle avoidance that complement the NETMORC robot controller developed in earlier years. The navigation/avoidance modules are able to modify the information about the angle to target in such a way that NETMORC will guide the robot around obstacles on its way to the target. Drawing again from biologically-motivated neural networks, Gaudiano and his collaborators have devised a *reactive* obstacle avoidance module, and a *predictive* obstacle avoidance module.

The reactive obstacle avoidance module is based on a phenomenon known as *peak shift*, which originates from the psychological literature on classical and instrumental conditioning. Mathematically, when a positive Gaussian (or any "bump-like" function) at one location is added to a negative Gaussian at a slightly shifted location, the resulting function exhibits a peak that is shifted from the peak of the positive Gaussian, *away* from the negative Gaussian.

The peak shift phenomenon can be used to perform navigation in a cluttered environment by representing the target as a positive Gaussian and the obstacles as negative Gaussians in a 2-D map of the environment. Such a map is shown in Figure 3(a). At each time step the module finds the peak of activity along a circular slice centered around the robot. The peak will represent the moment-by-moment direction that must be followed in order to reach the target while avoiding obstacles. This peak represents the moment-by-moment direction to be followed in order to approach the target without colliding with obstacles. The information is passed directly into the ANG population of NETMORC.

The reactive navigation module includes a simple dynamical scheme that prevents the module from becoming trapped in local minima even in the presence of concave obstacle configurations, which are especially problematic for potential function approaches: in this situation the robot would tend to travel around inside the concavity without being able to reach the target. The problem is solved by modulating the influence region of the obstacle Gaussians as a function of the robot's behavior. When the robot is not progressing to the target (e.g. backward movement) the widths of the obstacle Gaussians are increased, until the influence of the concave obstacle becomes wide enough to force the robot to avoid it altogether.

Figure 3 also shows performance under two conditions, one in which the robot must navigate in a maze-like environment, the other in which the robot must negotiate a concave obstacle directly in the way to the target. In the first case the robot reaches the target without trouble. In the second case, the robot first runs into the concave part of the obstacle, which forces it to turn away from the target. This generates a "frustration" signal that increases the width of the obstacle Gaussians, so that on the next pass the robot barely enters the concave obstacle, and finally navigates around it on the way to the goal.

Predictive Obstacle Avoidance Module

There are certain tasks for which reactive navigation is not sufficient. For instance, one might wish to modulate the influence of obstacles on the basis of factors other than proximity,

Figure 3. (a) Architecture of the reactive obstacle avoidance module, which combines target and obstacle information to extract a peak of activity along a circle around the robot. This peak of activity is used as the ANG information in the NETMORC architecture. Performance of the reactive obstacle avoidance module in the presence of (b) a maze-like environment, and (c) a concave obstacle blocking the way to the target.

13

Figure 4. (a) Conditioning model for obstacle avoidance. The ultrasonic range data represents the conditioned stimuli; the crash is the unconditioned stimulus. After conditioning, the pattern of activity across the ultrasonic sensors can predict a collision and change the angular velocity to avoid the obstacle. (b) Performance in the presence of several obstacles after the robot has been trained in an unsupervised fashion: the robot is able to follow a sequence of points, labeled 1-8, without colliding with the obstacles.

14

such as how important it is to reach the target. One example might be a robot that needs to reach a battery charging station: if the robot follows a long, roundabout path to avoid all obstacles, it may travel too far and run out of batteries. A related observation is that not all obstacles should be represented as equal sized Gaussians: avoiding a tree requires much less of a detour than avoiding a minefield or an unfriendly vehicle. In general, then, it would be desirable to be able to modulate the environment map on the basis of three forms of information: (1) sensory cues, e.g., the size of an obstacle; (2) internal cues, e.g., the need to return home or to explore novel objects; and (3) learned information about targets.

In last year's report, the PI proposed to develop a model for predictive obstacle avoidance based on Grossberg's model of classical and instrumental conditioning. Briefly, sensory cues from the environment are coupled to information about the organism's internal states, which in the case of the robot might include indications of battery level, distance from home, or internal damage. The circuit is designed so that when a strong noxious (or pleasant) stimulus occurs, such as a collision, the network learns the pattern of sensor activity that preceded the collision. Through higher-order conditioning, as the circuit experiences further collisions it learns to make earlier and earlier predictions on the basis of sensor inputs. How far back in time the prediction can be made depends on how quickly learning occurs: rapid learning means very few collisions, but it also means less high-order conditioning, and thus less anticipatory reaction to the collision.

The output of the conditioning circuit is used to modulate the sensor information as it projects to spatial maps of approach and avoidance. The combination of activity from these maps is equivalent to the difference-of-Gaussians map described in the previous section.

Professor Gaudiano and colleagues have implemented the above scheme in a simplified fashion by assuming that the robot has eight ultrasonic sensors whose activity is directly proportional to the distance to obstacles. The neural network model has no knowledge of the layout of the sensors on the robot surface: each sensor simply modulates the activity of one of the sensory nodes in the circuit. In addition, a bump sensor becomes active when a collision occurs (in practice, we assumed that a collision occurs when any one of the sensors sees an obstacle less than 1m away). The 2-D Gaussian map is simplified by projecting the sensor activations onto the 1-D angle map (ANG) of NETMORC, as shown in Figure 4(a). Finally, the output of the conditioning circuit learns to *inhibit* the pattern of angular velocities present on the angular map. By allowing each sensor to generate a Gaussian-like activity pattern on the angle map, the conditioning circuit learned to generate an inhibitory Gaussian-like pattern corresponding to the activity pattern that was present on the ANG population at the time of collision.

The network is trained by letting the robot make systematic movements at various combinations of wheel velocities in the presence of several obstacles. During learning, each sensor learned to predict what patterns of wheel velocities caused their activity to become very large (i.e., near an obstacle); for example, the sensor on the left side of the robot learned that left-going angular velocities led to collisions. After training, an obstacle on the left of the robot caused large activity in the left sensor, which projected an inhibitory Gaussian-like pattern to the left-going angular velocities. This inhibitory Gaussian interacted with the target Gaussian, leading to a shift in peak activity *toward the right* from the true target location. The performance of the module in the presence of a cluttered environment is shown in Figure 4(b). Notice that this scheme is based on *egocentric* coordinates, which means that learning is completely independent of the environment, and the robot is able to avoid obstacles in any configuration.

15

SELECTED ABSTRACTS

CEREBELLAR LEARNING IN AN OPPONENT MOTOR CONTROLLER FOR ADAPTIVE LOAD COMPENSATION AND SYNERGY FORMATION

Daniel Bullock†, José L. Contreras-Vidal‡, and Stephen Grossberg§

Abstract

This paper shows how a minimal neural network model of the cerebellum may be embedded within a sensory-neuro-muscular control system that mimics known anatomy and physiology. With this embedding, cerebellar learning promotes load compensation while also allowing both coactivation and reciprocal inhibition of sets of antagonist muscles. In particular, we show how synaptic long term depression guided by feedback from muscle stretch receptors can lead to trans-cerebellar gain changes that are load-compensating. It is argued that the same processes help to adaptively discover multi-joint synergies. Simulations of rapid single joint rotations under load illustrates design feasibility and stability.

17

A SELF-ORGANIZING NEURAL NETWORK FOR LEARNING A BODY-CENTERED INVARIANT REPRESENTATION OF 3-D TARGET POSITION

Daniel Bullock, Douglas Greve, Stephen Grossberg, and Frank H. Guenther

Abstract

This paper describes a self-organizing neural network that rapidly learns a body-centered representation of 3-D target positions. This representation remains invariant under head and eye movements, and is a key component of sensory-motor systems for producing motor equivalent reaches to targets (Bullock, Grossberg, and Guenther, 1993).

A SELF-ORGANIZING NEURAL MODEL OF MOTOR EQUIVALENT REACHING AND TOOL USE BY A MULTIJOINT ARM

Daniel Bullock†, Stephen Grossberg†, and Frank H. Guenther‡

Abstract

This paper describes a self-organizing neural model for eye-hand coordination. Called the DIRECT model, it embodies a solution of the classical motor equivalence problem. Motor equivalence computations allow humans and other animals to flexibly employ an arm with more degrees of freedom than the space in which it moves to carry out spatially defined tasks under conditions that may require novel joint configurations. During a motor babbling phase, the model endogenously generates movement commands that activate the correlated visual, spatial, and motor information that are used to learn its internal coordinate transformations. After learning occurs, the model is capable of controlling reaching movements of the arm to prescribed spatial targets using many different combinations of joints. When allowed visual feedback, the model can automatically perform, without additional learning, reaches with tools of variable lengths, with clamped joints, with distortions of visual input by a prism, and with unexpected perturbations. These compensatory computations occur within a single accurate reaching movement. No corrective movements are needed. Blind reaches using internal feedback have also been simulated. The model achieves its competence by transforming visual information about target position and end effector position in 3-D space into a body-centered spatial representation of the direction in 3-D space that the end effector must move to contact the target. The spatial direction vector is adaptively transformed into a motor direction vector, which represents the joint rotations that move the end effector in the desired spatial direction from the present arm configuration. Properties of the model are compared with psychophysical data on human reaching movements, neurophysiological data on the tuning curves of neurons in the monkey motor cortex, and alternative models of movement control.

19

A NEURAL NETWORK MODEL FOR CURSIVE SCRIPT PRODUCTION

Daniel Bullock†, Stephen Grossberg‡, and Christian Mannes§

Technical Report CAS/CNS TR-92-029
Boston, MA: Boston University
Biological Cybernetics, 1993, **70**, 15–28

Abstract

This article describes a neural network model, called the VITEWRITE model, for generating handwriting movements. The model consists of a sequential controller, or motor program, that interacts with a trajectory generator to move a hand with redundant degrees of freedom. The neural trajectory generator is the Vector Integration to Endpoint (VITE) model for synchronous variable-speed control of multijoint movements. VITE properties enable a simple control strategy to generate complex handwritten script if the hand model contains redundant degrees of freedom. The proposed controller launches transient directional commands to independent hand synergies at times when the hand begins to move, or when a velocity peak in a given synergy is achieved. The VITE model translates these temporally disjoint synergy commands into smooth curvilinear trajectories among temporally overlapping synergetic movements. The separate "score" of onset times used in most prior models is hereby replaced by a self-scaling activity-released "motor program" that uses few memory resources, enables each synergy to exhibit a unimodal velocity profile during any stroke, generates letters that are invariant under speed and size rescaling, and enables effortless connection of letter shapes into words. Speed and size rescaling are achieved by scalar GO and GRO signals that express computationally simple volitional commands. Psychophysical data concerning hand movements, such as the isochrony principle, asymmetric velocity profiles, and the two-thirds power law relating movement curvature and velocity arise as emergent properties of model interactions.

A SELF-ORGANIZING HEADING AND DEPTH DETECTION NETWORK

Seth Cameron†, Stephen Grossberg‡, and Frank H. Guenther§

Abstract

This paper describes a self-organizing neural network architecture that transforms optic flow information into representations of heading direction and distances to visible objects. The network's weights are trained during an action-perception cycle in which self-generated eye and body movements produce optic flow information. Learning of the relationship between eye movement outflow commands and rotational components of the flow field leads to a "translation only" flow field by removing the effects of eye movements from the original flow field. Categorization of normalized translation only flow patterns using a self-organizing feature map results in a map whose cells code heading directions. A third learning process combines this heading information with translation only flow field information to form a map whose cells each code the distance to the visible object (if any) at a particular retinal location. All learning processes take place concurrently and require no external "teachers". Simulations of the network verify its performance and its applicability to the problem of visual navigation.

MOTIVATION

Clark Dorman† and Paolo Gaudiano‡

Handbook of Brain Theory and Neural Networks

Abstract

The ability of humans and animals to survive in a constantly changing environment is a testament to the power of biological processes. At any moment in our lives, we are faced with many sensory stimuli, and we can typically generate a large number of behaviors. How do we learn to ignore irrelevant information and suppress inappropriate behavior so that we may function in a complex environment?

In this chapter we discuss *motivation*, the internal force that produces actions on the basis of the momentary balance between our needs and the demands of our environment. We first give a description of motivation and how it is studied, focusing on behavioral and physiological studies. We then discuss the role of motivation in behavioral theories and neural network modeling.

AN UNSUPERVISED ERROR-BASED NEURAL NETWORK MODEL FOR THE GENERATION AND CONTROL OF MOVEMENT TRAJECTORIES

Paolo Gaudiano†

Abstract

How can humans and animals be able to carry out novel motor tasks that they have never learned before? How is perceptual information about their environment transformed into spatial representations that can be used to generate accurate motor commands? In this talk I will present the Vector Associative Map (VAM), a self-organizing, unsupervised neural network model that has been applied to a variety of problems in the adaptive control of movemente trajectories. The VAM was derived from the Vector Integration To Endpoint (VITE) model (Bullock and Grossberg, 1988, *Psychological Review*, **95**, 49) for the generation and control of movement trajectories. The VAM model has been applied to a variety of learning tasks, including intramodal calibration of arm control parameters, intermodal learning of spatial-to-motor maps (Gaudiano and Grossberg, 1991, *Neural Networks*, **4**, 147), and learning an invariant representation of 3-D target positions in head-centered coordinates (Guenther, Bullock, Greve, Grossberg, *Journal of Cognitive Neuroscience*, in press). The VAM model advances our understanding of brain function in the realm of adaptive motor control, and it holds great potential for practical applications in robotics and control.

A NONLINEAR MODEL OF SPATIOTEMPORAL RETINAL PROCESSING: SIMULATIONS OF X AND Y RETINAL GANGLION CELL BEHAVIOR

Paolo Gaudiano†

Abstract

This article introduces a nonlinear model of neural processing in the vertebrate retina, comprising model photoreceptors, model push-pull bipolar cells, and model ganglion cells. Analyses and simulations show that the model can account for several aspects of both X and Y cat retinal ganglion cell behavior. In particular, with a choice of parameters that mimics *beta* cells, the model exhibits X-like linear spatial summation (null response to contrast-reversed gratings) in spite of photoreceptor nonlinearities; on the other hand, a choice of parameters that mimics *alpha* cells leads to Y-like frequency doubling. These and other results suggest that X and Y cells can be seen as variants of a single neural circuit. The model also suggests that both depolarizing and hyperpolarizing bipolar cells converge onto both ON and OFF ganglion cell types, although the effects of this push-pull convergence can be elusive when recording from individual ganglion cells. These hypotheses are supported in the article by a number of computer simulation results.

THE NEURAL DYNAMICS APPROACH TO SENSORY-MOTOR CONTROL: OVERVIEW AND RECENT APPLICATIONS IN MOBILE ROBOT CONTROL AND SPEECH PRODUCTION

Paolo Gaudiano†, Frank H. Guenther‡, and Eduardo Zalama

Invited article to appear in **Progress in Neural Networks**

Abstract

This chapter discusses a collection of models that utilize adaptive and dynamical properties of neural networks to solve problems of sensory-motor control for biological organisms and robots. The chapter begins with an overview of several unsupervised neural network models developed at the Center for Adaptive Systems during the past decade. These models have been used to explain a variety of data in research areas ranging from the cortical control of eye and arm movements to spinal regulation of muscle length and tension. Next, two recent models that build on important concepts from this earlier work are presented. The first of these models is an adaptive neural network controller for a visually-guided mobile robot. The neural network controller enables the robot to move to arbitrary targets without any knowledge of the robot's kinematics, immediately and automatically compensating for perturbations such as target movements, wheel slippage, or changes in the robot's plants. The controller also adapts to long-term perturbations, enabling the robot to compensate for statistically significant changes in its plant. The second model is a self-organizing neural network addressing speech motor skill acquisition and speech production. This model explains a wide range of data on contextual variability, motor equivalence, coarticulation, and speaking rate effects. Model parameters are learned during a babbling phase, using only information available to a babbling infant. After learning, the model can produce arbitrary phoneme strings, again exhibiting automatic compensation for perturbations or constraints on the articulators. Finally, other recent models using a neural dynamics approach are summarized and future research avenues are outlined.

A NEURAL NETWORK MODEL OF DYNAMIC RECEPTIVE FIELD REORGANIZATION

P. Gaudiano, S. Olson, D. Tal, and B. Fischl

Abstract

Primary sensory cortex is traditionally viewed as a passive filter, extracting information for processing in higher cortical centers. However, recent experiments have revealed a remarkable degree of plasticity in primary sensory cortex, particularly in visual cortex (Gilbert, 1992; Heinen & Skavenski, 1991; Kaas *et al.*, 1991) and somatosensory cortex (Merzenich *et al.*, 1984; Pons *et al.*, 1991; Ramachandran *et al.*, 1992). Receptive fields of cells in visual cortex have been shown to respond dynamically to changes in the visual environment, both within and outside the cells' classically defined receptive fields. This reorganization occurs on a variety of time scales, from seconds to years (Gilbert, 1992). We show a simple neural network model based on Adaptive Resonance Theory (ART: Carpenter & Grossberg, 1987; Grossberg, 1976) that displays some of the dynamical reorganization found in visual and somatosensory cortex. According to ART, plasticity is maintained throughout life, although feedback interactions prevent spurious reorganization during normal cortical function. In qualitative agreement with experimental results, simulated cortical cell receptive fields expand and contract as a result of attentional influences, real and artificial retinal lesions (both immediate and long-term reorganization), and preferential stimulation. Information from outside a cell's receptive field directly and indirectly mediates the cell's response. Both the rapid and long-term receptive field reorganizations arise as a consequence of nonlinear network-level interactions that are not fully explicable by examining the responses of individual neurons.

AN UNSUPERVISED NEURAL NETWORK FOR LOW-LEVEL CONTROL OF A WHEELED MOBILE ROBOT

P. Gaudiano†, E. Zalama, and J. López Coronado

Abstract

We have recently introduced a neural network mobile robot controller (NETMORC). This controller, based on previously developed neural network models of biological sensory-motor control, autonomously learns the forward and inverse odometry of a differential drive robot through an unsupervised learning-by-doing cycle. After an initial learning phase, the controller can move the robot to an arbitrary stationary or moving target while compensating for noise and other forms of disturbance, such as wheel slippage or changes in the robot's plant. In addition, the forward odometric map allows the robot to reach targets in the absence of sensory feedback. The controller is also able to adapt in response to long-term changes in the robot's plant, such as a change in the radius of the wheels. In this article we review the NETMORC architecture and describe its simplified algorithmic implementation, we present new, quantitative results on NETMORC's performance and adaptability under noise-free and noisy conditions, we compare NETMORC's performance on a trajectory-following task with the performance of an alternative controller, and we describe preliminary results on the hardware implementation of NETMORC with the mobile robot *ROBUTER*.

BOUNDARY, BRIGHTNESS, AND DEPTH INTERACTIONS DURING PREATTENTIVE REPRESENTATION AND ATTENTIVE RECOGNITION OF FIGURE AND GROUND

Stephen Grossberg†

Abstract

This article applies a recent theory of 3-D biological vision, called FACADE Theory, to explain several percepts which Kanizsa pioneered. These include 3-D pop-out of an occluding form in front of an occluded form, leading to completion and recognition of the occluded form; 3-D transparent and opaque percepts of Kanizsa squares, with and without Varin wedges; and interactions between percepts of illusory contours, brightness, and depth in response to 2-D Kanizsa images. These explanations clarify how a partially occluded object representation can be completed for purposes of object recognition, without the completed part of the representation necessarily being seen. The theory traces these percepts to neural mechanisms that compensate for measurement uncertainty and complementarity at individual cortical processing stages by using parallel and hierarchical interactions among several cortical processing stages. These interactions are modelled by a Boundary Contour System (BCS) that generates emergent boundary segmentations and a complementary Feature Contour System (FCS) that fills-in surface representations of brightness, color, and depth. The BCS and FCS interact reciprocally with an Object Recognition System (ORS) that binds BCS boundary and FCS surface representations into attentive object representations. The BCS models the parvocellular LGN→Interblob→Interstripe→V4 cortical processing stream, the FCS models the parvocellular LGN→Blob→Thin Stripe→V4 cortical processing stream, and the ORS models inferotemporal cortex.

SKILL ACQUISITION, COARTICULATION, AND RATE EFFECTS IN A NEURAL NETWORK MODEL OF SPEECH PRODUCTION

Frank H. Guenther†

Abstract

This work describes a neural network model of speech motor skill acquisition and speech production that explains a wide range of data on contextual variability, motor equivalence, coarticulation, and speaking rate effects. Model parameters are learned during a babbling phase. To explain how infants learn phoneme-specific and language-specific limits on acceptable articulatory variability, the learned speech sound targets take the form of regions, or convex hulls, in orosensory coordinates. This leads to an explanation of coarticulation wherein the target for a speech sound is reduced in size based on context to provide a more efficient sequence of articulator movements. Furthermore, reduction of target size for better accuracy during slower speech (in accordance with Fitt's law) leads to differential effects for vowels and consonants, as seen in speaking rate experiments that were previously explained by positing separate control processes for the two sound classes. The babbling process also naturally accounts for the formation of coordinative structures, or groups of articulator movements marshaled together to perform orosensory tasks. Coordinative structures provide motor equivalence, including automatic compensation to perturbations or constraints on the articulators. Computer simulations verify the model's motor equivalence, coarticulation, and speaking rate properties.

A NEURAL NETWORK MODEL OF SPEECH ACQUISITION AND MOTOR EQUIVALENT SPEECH PRODUCTION

Frank H. Guenther†

Abstract

This article describes a neural network model that addresses the acquisition of speaking skills by infants and subsequent motor equivalent production of speech sounds. The model learns two mappings during a babbling phase. An auditory-to-orosensory mapping specifies a vocal tract target for each speech sound; these targets take the form of convex hulls in orosensory coordinates defining the shape of the vocal tract. The babbling process wherein these convex hull targets are formed explains how an infant can learn phoneme-specific and language-specific limits on acceptable variability of articulator movements. The model also learns an orosensory-to-articulatory mapping wherein cells coding desired movement directions in orosensory space learn articulator movements that achieve these orosensory movement directions. The resulting mapping provides a natural explanation for the formation of coordinative structures. This mapping also makes efficient use of redundancy in the articulator system, thereby providing the model with motor equivalent capabilities. Simulations verify the model's ability to compensate for constraints or perturbations applied to the articulators automatically and without new learning and to explain contextual variability seen in human speech production.

SPEECH SOUND ACQUISITION, COARTICULATION, AND RATE EFFECTS IN A NEURAL NETWORK MODEL OF SPEECH PRODUCTION

Frank H. Guenther†

Abstract

This work describes a neural network model of speech motor skill acquisition and speech production that explains a wide range of data on contextual variability, motor equivalence, coarticulation, and speaking rate effects. Model parameters are learned during a babbling phase. To explain how infants learn phoneme-specific and language-specific limits on acceptable articulatory variability, the learned speech sound targets take the form of multidimensional regions, or convex hulls, in orosensory coordinates. Reduction of target size for better accuracy during slower speech (in the spirit of the speed-accuracy trade-off described by Fitts' law) leads to differential effects for vowels and consonants, as seen in speaking rate experiments that have been previously taken as evidence for separate control processes for the two sound types. An account of anticipatory coarticulation is posited wherein the target for a speech sound is reduced in size based on context to provide a more efficient sequence of articulator movements. This explanation generalizes the well-known look-ahead model of coarticulation to incorporate convex hull targets. Computer simulations verify the model's properties, including linear velocity/distance relationships, motor equivalence, speaking rate effects, and carryover and anticipatory coarticulation.

A MODELING FRAMEWORK FOR SPEECH MOTOR DEVELOPMENT AND KINEMATIC ARTICULATOR CONTROL

Frank H. Guenther†

Abstract

This paper presents three hypotheses that are central to a computational model of speech production: (1) Sound targets take the form of regions, rather than points, in a planning reference frame. (2) The planning frame is more acoustic-like than the frames used in most recent models. (3) A direction-to-direction mapping transforms planned trajectories into articulator movements. These hypotheses are supported by experimental data and simulation results.

A COMPUTATIONAL MODEL USING FORMANT SPACE PLANNING OF ARTICULATOR MOVEMENTS FOR VOWEL PRODUCTION

Frank H. Guenther† and Dave Johnson‡

Abstract

It is often hypothesized that articulator movements are planned within a coordinate frame whose variables correspond to key vocal tract constrictions [e.g., E. Saltzman and K.G. Munhall, *Ecological Psychology*, **1**, 333–382 (1989)]. However, recent evidence suggests that speakers may utilize a more acoustic-like space for planning vowel movements [J. Perkell, M. Matthies, M. Svirsky, and M. Jordan, *Journal of the Acoustical Society of America*, **93**, 2948–2961 (1993)]. Previous work has verified the capacity of a computational speech production model called DIVA to explain a wide range of experimental data using a constriction planning space. The current work extends the model to allow formant space planning of vowel movements. The model learns target regions for F1 and F2 for each vowel during a babbling cycle. A mapping between desired formant changes and articulator movements that achieve these changes is also learned. After babbling, the model successfully reaches all vowel targets from any initial vocal tract configuration, even in the presence of constraints such as a blocked jaw, and the resulting synthesized vowels are easily recognizable. Although vowel targets specify only formant ranges with no articulatory information, articulator configurations used by the model to produce vowels are similar to human configurations.

EFFICIENT CURVED REACHES RESULTING FROM KINEMATIC BIASES IN THE DIRECT MODEL

Frank H. Guenther† and Daniele Micci Barreca‡

In **Proceedings of the IEEE International Conference
on Systems, Man, and Cybernetics**
Vancouver, British Columbia, Canada, 1995

Abstract

The DIRECT model is a self-organizing neural network designed to explain neurophysi-ological and psychophysical data from targeted reaching experiments. The model's learning process is indirectly influenced by the arm's kinematics, resulting in movements biased toward joint rotations that produce the most spatial movement of the end-effector. This bias causes the end-effector trajectories performed after learning to deviate slightly from the straight path which would be produced by an explicit pseudoinverse computation, but the total joint rotation is significantly reduced by this slight curvature. A simplified model of this biasing is introduced, and implications regarding human arm movements are discussed.

34

ACOUSTIC SPACE MOVEMENT PLANNING IN A NEURAL MODEL OF MOTOR EQUIVALENT VOWEL PRODUCTION

Dave Johnson† and Frank H. Guenther‡

Abstract

Recent evidence suggests that speakers utilize an acoustic-like reference frame for the planning of speech movements. DIVA, a computational model of speech acquisition and motor equivalent speech production, has previously been shown to provide explanations for a wide range of speech production data using a constriction-based reference frame for movement planning. This paper extends the previous work by investigating an acoustic-like planning frame in the DIVA modeling framework. During a babbling phase, the model self-organizes targets in the planning space for each of ten vowels and learns a mapping from desired movement directions in this planning space into appropriate articulator velocities. Simulation results verify that after babbling the model is capable of producing easily recognizable vowel sounds using an acoustic planning space consisting of the formants F1 and F2. The model successfully reaches all vowel targets from any initial vocal tract configuration, even in the presence of constraints such as a blocked jaw.

EFFCIENT TRAJECTORY FORMATION
USING A LEARNED APPROXIMATE PSEUDOINVERSE
IN THE DIRECT MODEL OF REACHING

Daniele Micci-Barreca† and Frank H. Guenther‡

Abstract

The DIRECT model is a self-organizing neural network designed to explain neurophysiological and psychophysical data from targeted reaching experiments. A key transformation learned by the model is a mapping from spatial directions to joint rotations that approximates the pseudoinverse of the manipulator's Jacobian matrix. This learning occurs during an action-perception cycle and requires no external teacher and no explicit knowledge of the geometry of the manipulator. The learning process is indirectly influenced by the arm's kinematics, resulting in a bias toward joint rotations that produce the most spatial movement of the end-effector. This bias causes the end-effector trajectories performed after learning to deviate slightly from the straight path which would be produced by an explicit pseudoinverse computation, but the total joint rotation is significantly reduced by this slight curvature. The transformation learned by DIRECT therefore trades off absolute straightness of spatial trajectories for a substantial reduction of total joint movement. This property leads to an hypothesis concerning the coordinate frame used by humans to plan reaching movements.

36

NEURAL CONTROLLER FOR A MOBILE ROBOT
IN A STATIONARY ENVIRONMENT

F. Muñiz, E. Zalama, P. Gaudiano†, and J. López Coronado

Abstract

Recently we have introduced a neural controller for a mobile robot that learns both forward and inverse odometry of a differential-drive robot through an unsupervised learning-by-doing cycle. This article introduces an obstacle avoidance module that is integrated into the neural controller. This module makes use of sensory information to determine at each instant a desired angle and distance that causes the robot to navigate around obstacles on the way to a final target. Obstacle avoidance is performed in a reactive manner by representing the objects and target in the robot's environment as Gaussian functions. However, the influence of the Gaussians is modulated dynamically on the basis of the robot's behavior in a way that avoids problems with local minima. The proposed module enables the robot to operate successfully with different obstacle configurations, such as corridors, mazes, doors and even concave obstacles.

37

A COLLECTION OF ART-FAMILY GRAPHICAL SIMULATORS

David Pedini† and Paolo Gaudiano‡

Technical Report CAS/CNS-TR-94-023
Boston, MA: Boston University

Abstract

The Adaptive Resonance Theory (ART) architecture, first proposed by Grossberg (1976a, 1976b), is a self-organizing neural network for stable pattern categorization in response to arbitrary input sequences. Since its original formulation, several versions of ART have been proposed, each designed to handle a particular task or input format. Recent ART architectures have been designed to work in a supervised fashion, offering a viable alternative to supervised neural networks such as backpropagation (Rumelhart, 1986). Perhaps the best-known variant of ART is ART 2 (Carpenter and Grossberg, 1987), an unsupervised neural network that handles analog inputs. We have developed a series of simulators for some of the ART-family neural architectures, namely, ART 2 (Carpenter and Grossberg, 1987), ART2-A (Carpenter, Grossberg, and Rosen, 1991a), Fuzzy ART (Carpenter, Grossberg, and Rosen, 1991b), and Fuzzy ARTMAP (Carpenter, Grossberg, Markuzon, Reynolds, and Rosen, 1992). This article briefly summarizes the history and functionality of ART and its variants, and then describes the software package, which is available in the public domain.

A REAL-TIME, UNSUPERVISED NEURAL NETWORK MODEL FOR THE CONTROL OF A MOBILE ROBOT IN A NONSTATIONARY ENVIRONMENT

E. Zalama Casanova, Paolo Gaudiano†, and J. Lopéz Coronado

Abstract

This article introduces a real-time, unsupervised neural network model that learns to control a two-degree-of-freedom (2-DOF) nonholonomic mobile robot in a nonstationary environment. The model combines associative learning and Vector Associative Map (VAM) learning to generate transformations between spatial and velocity coordinates. The transformations are generated in an initial training phase, during which the robot moves as a result of endogenously generated velocities applied to the robot's wheels. The robot learns the relationship between these small velocities and the resulting incremental movements. During performance, the use of a VAM architecture enables the robot to generalize from the learned incremental movements to reach targets at arbitrary distance and angle from the robot. The VAM structure also enables the robot to perform successfully in spite of drastic changes to the robot's plant, including changes in wheel radius, changes in inter-wheel distance, or changes in the internal time step of the system. This article describes the model, presents illustrative simulation results that include both target and trajectory tracking, and compares the model to other neural network and classical approaches to control.

OBSTACLE AVOIDANCE BY MEANS OF
AN OPERANT CONDITIONING MODEL

E. Zalama, P. Gaudiano†, and J. López Coronado

Abstract

This paper describes the application of a model of operant conditioning to the problem of obstacle avoidance with a wheeled mobile robot. The main characteristic of the applied model is that the robot learns to avoid obstacles through a learning-by-doing cycle without external supervision. A series of ultrasonic sensors act as Conditioned Stimuli (CS), while collisions act as an Unconditioned Stimulus (UCS). By experiencing a series of movements in a cluttered environment, the robot learns to avoid sensor activation patterns that predict collisions, thereby learning to avoid obstacles. Learning generalizes to arbitrary cluttered environments. In this work we describe our initial implementation using a computer simulation.